

# Quantum Inference on Bayesian Networks

Guang Hao Low, Theodore J. Yoder, Isaac L. Chuang  
*Massachusetts Institute of Technology, 77 Massachusetts Avenue,  
Cambridge, 02139 MA, United States of America*  
(Dated: March 3, 2014)

Performing exact inference on Bayesian networks is known to be  $\#P$ -hard. Typically approximate inference techniques are used instead to sample from the distribution on query variables given the values  $e$  of evidence variables. Classically, a single unbiased sample is obtained from a Bayesian network on  $n$  variables with at most  $m$  parents per node in time  $\mathcal{O}(nmP(e)^{-1})$ , depending critically on  $P(e)$ , the probability the evidence might occur in the first place. By implementing a quantum version of rejection sampling, we obtain a square-root speedup, taking  $\mathcal{O}(n2^m P(e)^{-\frac{1}{2}})$  time per sample. We exploit the Bayesian network's graph structure to efficiently construct a quantum state, a q-sample, representing the intended classical distribution, and also to efficiently apply amplitude amplification, the source of our speedup. Thus, our speedup is notable as it is unrelativized – we count primitive operations and require no blackbox oracle queries.

PACS numbers: 02.50.Tt, 03.67.Ac

## I. INTRODUCTION

How are rational decisions made? Given a set of possible actions, the logical answer is the one with the largest corresponding utility. However, estimating these utilities accurately is the problem. A rational agent endowed with a model and partial information of the world must be able to evaluate the probabilities of various outcomes, and such is often done through inference on a Bayesian network [1], which efficiently encodes joint probability distributions in a directed acyclic graph of conditional probability tables. In fact, the standard model of a decision-making agent in a probabilistic time-discretized world, known as a Partially Observable Markov Decision Process, is a special case of a Bayesian network. Furthermore, Bayesian inference finds application in processes as diverse as system modeling [2], model learning [3, 4], data analysis [5], and decision making [6], all falling under the umbrella of machine learning [1].

Unfortunately, despite the vast space of applications, Bayesian inference is difficult. To begin with, exact inference is  $\#P$  hard in general [1]. It is often far more feasible to perform approximate inference by sampling, such as with the Metropolis-Hastings algorithm [7] and its innumerable specializations [8], but doing so is still NP-hard in general [9]. This can be understood by considering rejection sampling, a primitive operation common to many approximate algorithms that generates unbiased samples from a target distribution  $P(Q|\mathcal{E})$  for some set of query variables  $Q$  conditional on some assignment of evidence variables  $\mathcal{E} = e$ . In the general case, rejection sampling requires sampling from the full joint  $P(Q, \mathcal{E})$  and throwing away samples with incorrect evidence. In the specific case in which the joint distribution is described by a Bayesian network with  $n$  nodes each with no more than  $m$  parents, it takes time  $\mathcal{O}(nm)$  to generate a sample from the joint distribution, and so a sample from the conditional distribution  $P(Q|\mathcal{E})$  takes average time  $\mathcal{O}(nmP(e)^{-1})$ . Much of the computational difficulty is

related to how the marginal  $P(e) = P(\mathcal{E} = e)$  becomes exponentially small as the number of evidence variables increases, since only samples with the correct evidence assignments are recorded.

One very intriguing direction for speeding up approximate inference is in developing hardware implementations of sampling algorithms, for which promising results such as natively probabilistic computing with stochastic logic gates have been reported [10]. In this same vein, we could also consider physical systems that already describe probabilities and their evolution in a natural fashion to discover whether such systems would offer similar benefits.

Quantum mechanics can in fact describe such naturally probabilistic systems. Consider an analogy: if a quantum state is like a classical probability distribution, then measuring it should be analogous to sampling, and unitary operators should be analogous to stochastic updates. Though this analogy is qualitatively true and appealing, it is inexact in ways yet to be fully understood. Indeed, it is a widely held belief that quantum computers offer a strictly *more* powerful set of tools than classical computers, even probabilistic ones [11], though this appears difficult to prove [12]. Notable examples of the power of quantum computation include exponential speedups for finding prime factors with Shor's algorithm [13], and square-root speedups for generic classes of search problems through Grover's algorithm [14]. Unsurprisingly, there is an ongoing search for ever more problems amenable to quantum attack [15–17].

For instance, the quantum rejection sampling algorithm for approximate inference was only developed quite recently [18], alongside a proof, relativized by an oracle, of a square-root speedup in runtime over the classical algorithm. The algorithm, just like its classical counterpart, is an extremely general method of doing approximate inference, requiring preparation of a quantum pure state representing the joint distribution  $P(Q, \mathcal{E})$  and amplitude amplification to amplify the part of the superpo-

sition with the correct evidence. Owing to its generality, the procedure assumes access to a state-preparation oracle  $\hat{A}_P$ , and the runtime is therefore measured by the query complexity [19], the number of times the oracle must be used. Unsurprisingly, such oracles may not be efficiently implemented in general, as the ability to prepare arbitrary states allows for witness generation to QMA-complete problems [18, 20]. This also corresponds consistently to the NP-hardness of classical sampling.

In this paper, we present an unrelativized (*i.e.* no oracle) square-root, quantum speedup to rejection sampling on a Bayesian network. Just as the graphical structure of a Bayesian network speeds up classical sampling, we find that the same structure allows us to construct the state-preparation oracle  $\hat{A}_P$  efficiently. Specifically, quantum sampling from  $P(\mathcal{Q}|\mathcal{E} = e)$  takes time  $\mathcal{O}(n2^m P(e)^{-1/2})$ , compared with  $\mathcal{O}(nmP(e)^{-1})$  for classical sampling, where  $m$  is the maximum indegree of the network. We exploit the structure of the Bayesian network to construct an efficient quantum circuit  $\hat{A}_P$  composed of  $\mathcal{O}(n2^m)$  controlled-NOT gates and single-qubit rotations that generates the quantum state  $|\psi_P\rangle$  representing the joint  $P(\mathcal{Q}, \mathcal{E})$ . This state must then be evolved to  $|\mathcal{Q}\rangle$  representing  $P(\mathcal{Q}|\mathcal{E} = e)$ , which can be done by performing amplitude amplification [21], the source of our speedup and heart of quantum rejection sampling in general [18]. The desired sample is then obtained in a single measurement of  $|\mathcal{Q}\rangle$ .

We better define the problem of approximate inference with a review of Bayesian networks in section II. We discuss the sensible encoding of a probability distribution in a quantum state axiomatically in section III. This is followed by an overview of amplitude amplification in section IV. The quantum rejection sampling algorithm is given in section V. As our main result, we construct circuits for the state preparation operator in sections VIA and VIB and circuits for the reflection operators for amplitude amplification in section VIC. The total time complexity of quantum rejection sampling in Bayesian networks is evaluated in section VID, and we present avenues for further work in section VII.

## II. BAYESIAN NETWORKS

A Bayesian network is a directed acyclic graph structure that represents a joint probability distribution over  $n$  bits. A significant advantage of the Bayesian network representation is that the space complexity of the representation can be made much less than the general case by exploiting conditional dependencies in the distribution. This is achieved by associating with each graph node a conditional probability table for each random variable, with directed edges representing conditional dependencies, such as in Fig. 1a.

We adopt the standard convention of capital letters (e.g.  $X$ ) representing random *variables* while lowercase letters (e.g.  $a$ ) are particular fixed *values* of those vari-

ables. For simplicity, the random variables are taken to be binary. Accordingly, probability vectors are denoted  $P(X) = \{P(X=0), P(X=1)\}$  while  $P(x) \equiv P(X=x)$ . Script letters represent a set of random variables  $\mathcal{X} = \{X_1, X_2, \dots, X_n\}$ .

An arbitrary joint probability distribution  $P(x_1, x_2, \dots, x_n)$  on  $n$  bits can always be factored by recursive application of Bayes' rule  $P(X, Y) = P(X)P(Y|X)$ ,

$$P(x_1, x_2, \dots, x_n) = P(x_1) \prod_{i=2}^n P(x_i | x_1, \dots, x_{i-1}). \quad (1)$$

However, in most practical situations a given variable  $X_i$  will be dependent on only a few of its predecessors' values, those we denote by  $\text{parents}(X_i) \subseteq \{x_1, \dots, x_{i-1}\}$  (see Fig. 1a). Therefore, the factorization above can be simplified to

$$P(x_1, x_2, \dots, x_n) = P(x_1) \prod_{i=2}^n P(x_i | \text{parents}(X_i)). \quad (2)$$

A Bayes net diagrammatically expresses this simplification, with a topological ordering on the nodes  $X_1 \preceq X_2 \preceq \dots \preceq X_n$  in which parents are listed before their children. With a node  $x_i$  in the Bayes net, the conditional probability factor  $P(x_i = 1 | \text{parents}(X_i))$  is stored as a table of  $2^{m_i}$  values [1] where  $m_i$  is the number of parents of node  $X_i$ , also known as the indegree. Letting  $m$  denote the largest  $m_i$ , the Bayes net data structure stores at most  $\mathcal{O}(n2^m)$  probabilities, a significant improvement over the direct approach of storing  $\mathcal{O}(2^n)$  probabilities [1].

A common problem with any probability distribution is inference. Say we have a complete joint probability distribution on  $n$ -bits,  $P(\mathcal{X})$ . Given the values  $e = e_{|\mathcal{E}|} \dots e_2 e_1$  for a set  $\mathcal{E} \subseteq \mathcal{X}$  of random variables, the task is to find the distribution over a collection of query variables  $\mathcal{Q} \subseteq \mathcal{X} \setminus \mathcal{E}$ . That is, the exact inference problem is to calculate  $P(\mathcal{Q}|\mathcal{E} = e)$ . Exact inference is #P-hard [1], since one can create a Bayes net encoding the  $n$  variable  $k$ -SAT problem, with nodes for each variable, each clause, and the final verdict – a count of the satisfying assignments.

Approximate inference on a Bayesian network is much simpler, thanks to the graphical structure. The procedure for sampling is as follows: working from the top of the network, generate a value for each node, given the values already generated for its parents. Since each node has at most  $m$  parents that we must inspect before generating a value, and there are  $n$  nodes in the tree, obtaining a sample  $\{x_1, x_2, \dots, x_n\}$  takes time  $\mathcal{O}(nm)$ . Yet we must postselect on the correct evidence values  $\mathcal{E} = e$ , leaving us with an average time per sample of  $\mathcal{O}(nmP(e)^{-1})$ , which suffers when the probability  $P(e)$  becomes small, typically exponentially small with the number of evidence variables  $|\mathcal{E}|$ . Quantum rejection sampling, however, will improve the factor of  $P(e)^{-1}$  to  $P(e)^{-1/2}$ , while preserving the linear scaling in the number of variables  $n$ , given

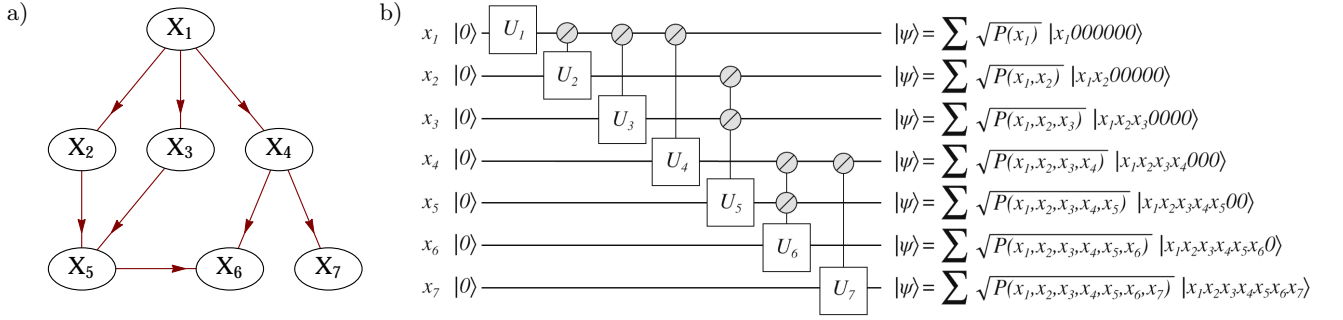


FIG. 1: a) An example of a directed acyclic graph which can represent a Bayesian network by associating with each node a conditional probability table. For instance, associated with the node  $X_1$  is the one value  $P(X_1 = 1)$ , while that of  $X_5$  consists of four values, the probabilities  $X_5 = 1$  given each setting of the parent nodes,  $X_2$  and  $X_3$ . b) A quantum circuit that efficiently prepares the q-sample representing the full joint distribution of (a). Notice in particular how the edges in the graph are mapped to conditioning nodes in the circuit. The  $|\psi_j\rangle$  represent the state of the system after applying the operator sequence  $U_1 \dots U_j$  to the initial state  $|0000000\rangle$ .

that we use an appropriate quantum state to represent the Bayesian network.

### III. QUANTUM SAMPLING FROM $P(\mathcal{X})$

This section explores the analogy between quantum states and classical probability distributions from first principles. In particular, for a classical probability distribution function  $P(\mathcal{X})$  on a set of  $n$  binary random variables  $\mathcal{X}$  what quantum state  $\rho_P$  (possibly mixed,  $d$ -qubits) should we use to represent it? The suitable state, which we call a quantum probability distribution function (qpdf), is defined with three properties.

**Definition 1.** A qpdf for the probability distribution  $P(\mathcal{X})$  has the following three properties:

1. Purity: In the interest of implementing quantum algorithms, we require the qpdf be a pure state  $\rho_P = |\Psi_P\rangle\langle\Psi_P|$ .
2. Q-sampling: A single qpdf can be measured to obtain a classical  $n$ -bit string, a sample from  $P(\mathcal{X})$ . Furthermore, for any subset of variables  $\mathcal{W} \subset \mathcal{X}$ , a subset of qubits in the qpdf can be measured to obtain a sample from the marginal distribution  $P(\mathcal{W})$ . We call these measurement procedures q-sampling.
3. Q-stochasticity: For every stochastic matrix  $T$  there is a unitary  $U_T$  such that whenever  $T$  maps the classical distribution  $P(\mathcal{X})$  to  $P'(\mathcal{X})$ ,  $U_T$  maps the qpdf  $|\Psi_P\rangle$  to  $|\Psi_{P'}\rangle = U_T |\Psi_P\rangle$ .

The motivation for property 3 is for implementing Markov chains, Markov decision processes, or even sampling algorithms such as Metropolis-Hastings, on quantum states. The question we pose, and leave open, is whether a qpdf exists.

The simplest way to satisfy the first two criteria, but not the third, is to initialize a single qubit for each classical binary random variable. This leads to what is called the q-sample, defined in prior work [22] as:

**Definition 2.** The q-sample of the joint distribution  $P(x_1, \dots, x_n)$  over  $n$  binary variables  $\{X_i\}$  is the  $n$ -qubit pure state  $|\psi_P\rangle = \sum_{x_1, \dots, x_n} \sqrt{P(x_1, \dots, x_n)} |x_1 \dots x_n\rangle$ .

The q-sample possesses property 1 and the eponymous property 2 above. However, it does not allow for stochastic updates as per property 3, as a simple single qubit example shows. In that case, property 3 requires

$$\begin{pmatrix} U_{11} & U_{12} \\ U_{21} & U_{22} \end{pmatrix} \begin{pmatrix} \sqrt{p} \\ \sqrt{1-p} \end{pmatrix} = \begin{pmatrix} \sqrt{pT_{11} + (1-p)T_{12}} \\ \sqrt{pT_{21} + (1-p)T_{22}} \end{pmatrix}, \quad (3)$$

for all  $p \in [0, 1]$ . Looking at Eq. 3 for  $p = 0$  and  $p = 1$  constrains  $U$  completely, and it is never unitary when  $T$  is stochastic. Thus, the q-sample fails to satisfy property 3.

Yet, the q-sample satisfies properties 1 and 2 in a very simplistic fashion, and various more complicated forms might be considered. For instance, relative phases could be added to the q-sample giving  $\sum_x e^{i\phi(x)} \sqrt{P(x)} |x\rangle$ , though this alone does not guarantee property 3, which is easily checked by adding phases to the proof above. Other extensions of the q-sample may include ancilla qubits, different measurement bases, or a post-processing step including classical randomness to translate the measurement result into a classical sample. It is an open question whether a more complicated representation satisfying all three properties exists, including q-stochasticity, so that we would have a qpdf possessing all the defining properties.

Nevertheless, although the q-sample is not a qpdf by our criteria, it will still be very useful for performing quantum rejection sampling. The property that a sample from a marginal distribution is obtained by simply

measuring a subset of qubits means that, using conditional gates, we can form a q-sample for a conditional distribution from the q-sample for the full joint distribution, as we will show in section VI. This corresponds to the classical formula  $P(\mathcal{V}|\mathcal{W}) = P(\mathcal{V}, \mathcal{W})/P(\mathcal{W})$ , which is the basis behind rejection sampling. The way it is actually done quickly on a q-sample is through amplitude amplification, reviewed next, in the general case.

#### IV. AMPLITUDE AMPLIFICATION

Amplitude amplification [21] is a well-known extension of Grover's algorithm and is the second major concept in the quantum inference algorithm. Given a quantum circuit  $\hat{A}$  for the creation of an  $n$ -qubit pure state  $|\psi\rangle = \alpha|\psi_t\rangle + \beta|\bar{\psi}_t\rangle = \hat{A}|0\rangle^{\otimes n}$ , where  $\langle\psi_t|\bar{\psi}_t\rangle = 0$ , the goal is to return the target state  $|\psi_t\rangle$  with high probability. To make our circuit constructions more explicit, we assume target states are marked by a known evidence bit string  $e = e_{|\mathcal{E}|}\dots e_2e_1$ , so that  $|\psi_t\rangle = |\mathcal{Q}\rangle|e\rangle$  lives in the tensor product space  $\mathcal{H}_{\mathcal{Q}} \otimes \mathcal{H}_{\mathcal{E}}$  and the goal is to extract  $|\mathcal{Q}\rangle$ .

Just like in Grover's algorithm, a pair of reflection operators are applied repetitively to rotate  $|\psi\rangle$  into  $|\psi_t\rangle$ . Reflection about the evidence is performed by  $\hat{S}_e = \hat{I} \otimes (\hat{I} - 2|e\rangle\langle e|)$  followed by reflection about the initial state by  $\hat{S}_\psi = (\hat{I} - 2|\psi\rangle\langle\psi|)$ . Given  $\hat{A}$ , then  $\hat{S}_\psi = \hat{A}\hat{S}_0\hat{A}^\dagger$ , where  $\hat{S}_0 = (\hat{I} - 2|0\rangle\langle 0|^{\otimes n})$ .

The analysis of the amplitude amplification algorithm is elucidated by writing the Grover iterate  $\hat{G} = -\hat{S}_\psi\hat{S}_e = -\hat{A}\hat{S}_0\hat{A}^\dagger\hat{S}_e$  in the basis of  $\frac{\alpha}{|\alpha|}|\psi_t\rangle \equiv \begin{pmatrix} 1 \\ 0 \end{pmatrix}$  and  $\frac{\beta}{|\beta|}|\bar{\psi}_t\rangle \equiv \begin{pmatrix} 0 \\ 1 \end{pmatrix}$  [19],

$$\hat{G} = \begin{pmatrix} 1 - 2|\alpha|^2 & 2|\alpha|\sqrt{1 - |\alpha|^2} \\ -2|\alpha|\sqrt{1 - |\alpha|^2} & 1 - 2|\alpha|^2 \end{pmatrix}. \quad (4)$$

In this basis, the Grover iterate corresponds to a rotation by small angle  $\theta = \cos^{-1}(1 - 2|\alpha|^2) \approx 2|\alpha|$ . Therefore, applying the iterate  $N$  times rotates the state by  $N\theta$ . We conclude that  $\hat{G}^N|\psi\rangle$  is closest to  $\frac{\alpha}{|\alpha|}|\psi_t\rangle$  after  $N = \mathcal{O}\left(\frac{\pi}{4|\alpha|}\right)$  iterations.

Usually, amplitude amplification needs to be used without knowing the value of  $|\alpha|$ . In that case,  $N$  is not known. However, the situation is remedied by guessing the correct number of Grover iterates to apply in exponential progression. That is, we apply  $\hat{G}$   $2^k$  times, with  $k = 0, 1, 2, \dots$ , measure the evidence qubits  $|\mathcal{E}\rangle$  after each attempt, and stop when we find  $\mathcal{E} = e$ . It has been shown [21] that this approach also requires on average  $\mathcal{O}\left(\frac{1}{|\alpha|}\right)$  applications of  $\hat{G}$ .

#### V. THE QUANTUM REJECTION SAMPLING ALGORITHM

The quantum rejection sampling algorithm [18], which we review now, is an application of amplitude amplification on a q-sample. The general problem, as detailed in section II, is to sample from the  $n$ -bit distribution  $P(\mathcal{Q}|\mathcal{E} = e)$ . We assume that we have a circuit  $\hat{A}_P$  that can prepare the q-sample  $|\psi_P\rangle = \hat{A}_P|0\rangle^{\otimes n}$ . Now, permuting qubits so the evidence lies to the right, the q-sample can be decomposed into a superposition of states with correct evidence and states with incorrect evidence.

$$|\psi_P\rangle = \sqrt{P(e)}|\mathcal{Q}\rangle|e\rangle + \sqrt{1 - P(e)}|\bar{\mathcal{Q}}, e\rangle, \quad (5)$$

where  $|\mathcal{Q}\rangle$  denotes the q-sample of  $P(\mathcal{Q}|\mathcal{E} = e)$ , our target state. Next perform the amplitude amplification algorithm from the last section to obtain  $|\mathcal{Q}\rangle$  with high probability. Note that this means the state preparation operator  $\hat{A}_P$  must be applied  $\mathcal{O}(P(e)^{-1/2})$  times. Once obtained,  $|\mathcal{Q}\rangle$  can be measured to get a sample from  $P(\mathcal{Q}|\mathcal{E} = e)$ , and we have therefore done approximate inference. Pseudocode is provided as an algorithm 1.

However, we are so far missing a crucial element. How is the q-sample preparation circuit  $\hat{A}_P$  actually implemented, and can this implementation be made efficient, that is, polynomial in the number of qubits  $n$ ? The answer to this question removes the image of  $\hat{A}_P$  as a featureless black box and is addressed in the next section.

---

**Algorithm 1** Quantum rejection sampling algorithm: generate one sample from  $P(\mathcal{Q}|\mathcal{E} = e)$  given a q-sample preparation circuit  $\hat{A}_P$

---

```

 $k \leftarrow -1$ 
while evidence  $\mathcal{E} \neq e$  do
   $k \leftarrow k + 1$ 
   $|\psi_P\rangle \leftarrow \hat{A}_P|0\rangle^{\otimes n}$  //prepare a q-sample of  $P(\mathcal{X})$ 
   $|\psi'_P\rangle \leftarrow \hat{G}^{2^k}|\psi_P\rangle$  //where  $\hat{G} = -\hat{A}_P\hat{S}_0\hat{A}_P^\dagger\hat{S}_e$ 
  Measure evidence qubits  $\mathcal{E}$  of  $|\psi'_P\rangle$ 
  Measure the query qubits to obtain a sample  $\mathcal{Q} = q$ 

```

---

#### VI. CIRCUIT CONSTRUCTIONS

While the rejection sampling algorithm from Section V is entirely general for any distribution  $P(\mathcal{X})$ , the complexity of q-sample preparation, in terms of the total number of CNOTs and single qubit rotations involved, is generally exponential in the number of qubits,  $\mathcal{O}(2^n)$ . We show this in section VIA. The difficulty is not surprising, since arbitrary q-sample preparation encompasses witness generation to QMA-complete problems [18, 20]. However, there are cases in which the q-sample can be prepared efficiently [22]. The main result of this paper is that, for probability distributions resulting from a Bayesian network  $\mathcal{B}$  with  $n$  nodes and maximum indegree  $m$ , the circuit complexity of the q-sample preparation circuit  $\hat{A}_{\mathcal{B}}$  is  $\mathcal{O}(n2^m)$ . We show this in section

VIB. The circuit constructions for the remaining parts of the Grover iterate, the phase flip operators, are given in section VIC. Finally, we evaluate the complexity of our constructions as a whole in section VID and find that approximate inference on Bayesian networks can be done with a polynomially sized quantum circuit.

Throughout this section we will denote the circuit complexity of a circuit  $\hat{C}$  as  $Q_{\hat{C}}$ . This complexity measure is the count of the number of gates in  $\hat{C}$  after compilation into a complete, primitive set. The primitive set we employ includes the CNOT gate and all single qubit rotations.

### A. Q-sample Preparation

If  $P(x)$  lacks any kind of structure, the difficulty of preparing the q-sample  $|\psi_P\rangle = \hat{A}_P |0\rangle^{\otimes n}$  with some unitary  $\hat{A}_P$  scales at least exponentially with the number of qubits  $n$  in the q-sample. Since  $P(x)$  contains  $2^n - 1$  arbitrary probabilities,  $\hat{A}_P$  must contain at least that many primitive operations. In fact, the bound is tight — we can construct a quantum circuit preparing  $|\psi_P\rangle$  with complexity  $\mathcal{O}(2^n)$ .

**Theorem 1.** *Given an arbitrary joint probability distribution  $P(x_1, \dots, x_n)$  over  $n$  binary variables  $\{X_i\}$ , there exists a quantum circuit  $\hat{A}_P$  that prepares the q-sample  $\hat{A}_P |0\rangle^{\otimes n} = |\psi_P\rangle = \sum_{x_1, \dots, x_n} \sqrt{P(x_1, \dots, x_n)} |x_1 \dots x_n\rangle$  with circuit complexity  $\mathcal{O}(2^n)$ .*

*Proof.* Decompose  $P(x) = P(x_1) \prod_{i=2}^n P(x_i | x_1 \dots x_{i-1})$  as per Eq. (1). For each conditional distribution  $P(X_i | x_1 \dots x_{i-1})$ , let us define the  $i$ -qubit uniformly controlled rotation  $\hat{U}_i$  such that given an  $(i-1)$  bit string assignment  $x_c \equiv x_1 \dots x_{i-1}$  on the control qubits, the action of  $\hat{U}_i$  on the  $i^{\text{th}}$  qubit initialized to  $|0\rangle_i$  is a rotation about the  $y$ -axis by angle  $2 \tan^{-1}(\sqrt{P(x_i=1|x_c)/P(x_i=0|x_c)})$  or  $\hat{U}_i |0\rangle_i = \sqrt{P(x_i=0|x_c)} |0\rangle_i + \sqrt{P(x_i=1|x_c)} |1\rangle_i$ . With this definition, the action of the single-qubit  $\hat{U}_1$  is  $\hat{U}_1 |0\rangle_1 = \sqrt{P(x_1=0)} |0\rangle_1 + \sqrt{P(x_1=1)} |1\rangle_1$ . By applying Bayes' rule in reverse, the operation  $\hat{A}_P = \hat{U}_n \dots \hat{U}_1$  then produces  $|\psi_P\rangle = \hat{A}_P |0\rangle$ . As each  $k$ -qubit uniformly controlled rotation is decomposable into  $\mathcal{O}(2^k)$  CNOTs and single-qubit rotations [23], the circuit complexity of  $\hat{A}_P$  is  $Q_{\hat{A}_P} = \sum_{i=1}^n \mathcal{O}(2^{i-1}) = \mathcal{O}(2^n)$ .  $\square$

The key quantum compiling result used in this proof is the construction of Bergholm et. al. [23] that decomposes  $k$ -qubit uniformly controlled gates into  $\mathcal{O}(2^k)$  CNOTs and single qubit operations. Each uniformly controlled gate is the realization of a conditional probability table from the factorization of the joint distribution. We use this result again in Bayesian q-sample preparation.

### B. Bayesian Q-sample Preparation

We now give our main result, a demonstration that the circuit  $\hat{A}_{\mathcal{B}}$  that prepares the q-sample of a Bayesian network is exponentially simpler than the general q-sample preparation circuit  $\hat{A}_P$ . We begin with a Bayesian network with, as usual,  $n$  nodes and maximum indegree  $m$  that encodes a distribution  $P(\mathcal{X})$ . As a minor issue, because the Bayesian network may have nodes reordered, the indegree  $m$  is actually a function of the specific parentage of nodes in the tree. This non-uniqueness of  $m$  corresponds to the non-uniqueness of the decomposition  $P(x_1, \dots, x_n) = P(x_1) \prod_{i=2}^n P(x_i | x_1 \dots x_{i-1})$  due to permutations of the variables. Finding the variable ordering minimizing  $m$  is unfortunately an NP-hard problem [24], but typically the variables have real-world meaning and the natural causal ordering often comes close to optimal [25]. In any case, we take  $m$  as a constant much less than  $n$ .

**Definition 3.** If  $P(\mathcal{X})$  is the probability distribution represented by a Bayesian network  $\mathcal{B}$ , the Bayesian q-sample  $|\psi_{\mathcal{B}}\rangle$  denotes the q-sample of  $P(\mathcal{X})$ .

**Theorem 2.** *The Bayesian q-sample of the Bayesian network  $\mathcal{B}$  with  $n$  nodes and bounded indegree  $m$  can be prepared efficiently by an operator  $\hat{A}_{\mathcal{B}}$  with circuit complexity  $\mathcal{O}(n2^m)$  acting on the initial state  $|0\rangle^{\otimes n}$ .*

*Proof.* As a Bayesian network is a directed acyclic graph, let us order the node indices topologically such that for all  $1 \leq i \leq n$ , we have  $\text{parents}(x_i) \subseteq \{x_1, x_2, \dots, x_{i-1}\}$ , and  $\max_i |\text{parents}(x_i)| = m$ . Referring to the construction from the proof of theorem 1, the state preparation operator  $\hat{A} = \hat{U}_n \dots \hat{U}_1$  then contains at most  $m$ -qubit uniformly controlled operators, each with circuit complexity  $\mathcal{O}(2^m)$ , again from Bergholm et. al. [23]. The circuit complexity of  $\hat{A}_{\mathcal{B}}$  is thus  $Q_{\hat{A}_{\mathcal{B}}} = \sum_{i=1}^n \mathcal{O}(2^m) = \mathcal{O}(n2^m)$ .  $\square$

Fig. 1b shows the circuit we have just described. Bayesian q-sample preparation forms part of the Grover iterate required for amplitude amplification. The rest is comprised of the reflection, or phase flip, operators.

### C. Phase Flip Operators

Here we show that the phase flip operators are also efficiently implementable, so that we can complete the argument that amplitude amplification on a Bayesian q-sample is polynomial time. Note first that the phase flip operators  $\hat{S}_e$  acting on  $k = |\mathcal{E}| \leq n$  qubits can be implemented with a single  $k$ -qubit controlled  $Z$  operation along with at most  $2k$  bit flips. The operator  $\hat{S}_0$  is the special case  $\hat{S}_{e=0^n}$ . A Bayesian q-sample can be decomposed exactly as in Eq. (5)

$$|\psi_{\mathcal{B}}\rangle = \sqrt{P(e)} |\mathcal{Q}\rangle |e\rangle + \sqrt{1 - P(e)} |\overline{\mathcal{Q}}, e\rangle. \quad (6)$$

Recall  $|\mathcal{Q}\rangle$  is the q-sample of  $P(\mathcal{Q}|\mathcal{E} = e)$  and  $|\overline{\mathcal{Q}e}\rangle$  contains all states with invalid evidence  $\mathcal{E} \neq e$ . We write the evidence as a  $k$ -bit string  $e = e_k \dots e_2 e_1$  and  $\hat{X}_i$  as the bit flip on the  $i^{\text{th}}$  evidence qubit. The controlled phase, denoted  $\hat{Z}_{1\dots k}$ , acts on all  $k$  evidence qubits symmetrically, flipping the phase if and only if all qubits are 1. Then  $\hat{S}_e$  is implemented by

$$\hat{S}_e = \hat{B} \hat{Z}_{1\dots k} \hat{B} \quad (7)$$

where  $\hat{B} = \prod_{i=1}^k \hat{X}_i^{\bar{e}_i}$  with  $\bar{e}_i \equiv 1 - e_i$ . Explicitly,

$$\begin{aligned} \hat{S}_e |\psi_{\mathcal{B}}\rangle &= \hat{B} \hat{Z}_{1\dots k} \hat{B} \left[ \sqrt{P(e)} |\mathcal{Q}\rangle |e\rangle + \sqrt{1 - P(e)} |\overline{\mathcal{Q}e}\rangle \right] \\ &= \hat{B} \hat{Z}_{1\dots k} \left[ \sqrt{P(e)} |\mathcal{Q}\rangle |1^n\rangle + \sqrt{1 - P(e)} |\overline{\mathcal{Q}1^n}\rangle \right] \\ &= \hat{B} \left[ -\sqrt{P(e)} |\mathcal{Q}\rangle |1^n\rangle + \sqrt{1 - P(e)} |\overline{\mathcal{Q}1^n}\rangle \right] \\ &= \left[ -\sqrt{P(e)} |\mathcal{Q}\rangle |e\rangle + \sqrt{1 - P(e)} |\overline{\mathcal{Q}e}\rangle \right]. \end{aligned} \quad (8)$$

The circuit diagram representing  $\hat{S}_e$  is shown in Fig. 2. The  $k$ -qubit controlled phase can be constructed from  $\mathcal{O}(k)$  CNOTs and single qubit operators using  $\mathcal{O}(k)$  ancillas [16] or, alternatively,  $\mathcal{O}(k^2)$  CNOTs and single qubit operators using no ancillas [26].

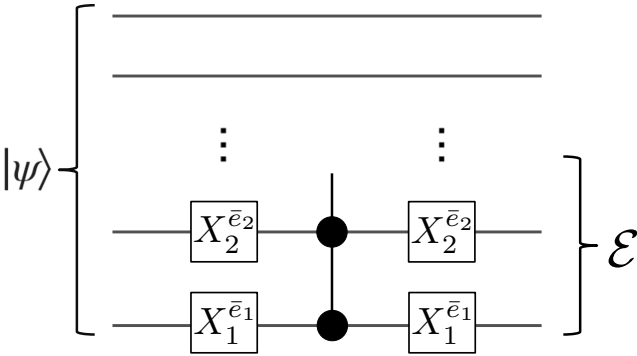


FIG. 2: Quantum circuit for implementing the phase flip operator  $\hat{S}_e$ . The  $k = |\mathcal{E}|$  qubit controlled phase operator acts on the evidence qubits  $\mathcal{E}$ . It may be compiled into  $\mathcal{O}(k)$  CNOTs and single qubit operators given  $\mathcal{O}(k)$  ancillas [16]. The evidence values  $e = e_k \dots e_2 e_1$  control the bit flips through  $\bar{e}_i \equiv 1 - e_i$ .

$\hat{U}$	$Q_{\hat{U}}$	Comments
$\hat{A}_P$	$\mathcal{O}(2^n)$	Q-sample preparation
$\hat{A}_B$	$\mathcal{O}(n2^m)$	Bayesian state preparation
$\hat{S}_0$	$\mathcal{O}(n)$	$\mathcal{O}(n)$ ancilla qubits
$\hat{S}_e$	$\mathcal{O}( \mathcal{E} )$	$\mathcal{O}( \mathcal{E} )$ ancilla qubits

TABLE I: Circuit complexity  $Q_{\hat{U}}$  of implementing the operators  $\hat{U}$  discussed in the text. The Grover iterate  $\hat{G}$  for amplitude amplification of a Bayesian q-sample (general q-sample) consists of two instances of the preparation circuit  $\hat{A}_B$  ( $\hat{A}_P$ ) and one instance each of  $\hat{S}_0$  and  $\hat{S}_e$ . The time to collect one sample from  $P(\mathcal{Q}|\mathcal{E} = e)$  is  $\mathcal{O}(Q_{\hat{G}}P(e)^{-1/2})$ .

#### D. Time Complexity

The circuit complexities of the various elements in the Grover iterate  $\hat{G} = -\hat{A}\hat{S}_0\hat{A}^\dagger\hat{S}_e$  are presented in Table I. As the circuit complexity of the phase flip operator  $S_0$  ( $S_e$ ) scales linearly with number of qubits  $n$  ( $|\mathcal{E}|$ ),  $Q_{\hat{G}}$  is dominated by the that of the state preparation operator  $\hat{A}$ . Although  $Q_{\hat{A}_P}$  scales exponentially with the number of nodes  $n$  for general q-sample preparation, Bayesian q-sample preparation on a network of bounded indegree  $m$  is efficient. Namely,  $Q_{\hat{A}_B} = \mathcal{O}(n2^m)$  scales linearly with  $n$  as in classical sampling from a Bayesian network. It takes  $\mathcal{O}(P(e)^{-1/2})$  applications of  $\hat{A}_B$  to perform the rejection sampling algorithm from Section V and, thus, a single sample from  $P(\mathcal{Q}|\mathcal{E})$  can be obtained by a quantum computer in time  $\mathcal{O}(n2^mP(e)^{-1/2})$ . In Section II, we saw that classical Bayesian inference takes time  $\mathcal{O}(nmP(e)^{-1})$  to generate a single sample. Thus, quantum inference on a Bayesian network provides a square-root speedup over the classical case. The quantum circuit diagram for Bayesian inference is outlined in Fig. 3.

#### VII. CONCLUSION

We have shown how the structure of a Bayesian network allows for a square-root, quantum speedup in approximate inference. We explicitly constructed a quantum circuit from CNOT and single qubit rotations that returns a sample from  $P(\mathcal{Q}|\mathcal{E} = e)$  using just  $\mathcal{O}(n2^mP(e)^{-1/2})$  gates. For more general probability distributions, the Grover iterate would include a quantity of gates exponential in  $n$ , the number of random variables, and thus not be efficient. This efficiency of our algorithm implies experimental possibilities. As a proof of principle, one could experimentally perform inference on a two node Bayesian network with only two qubits with current capabilities of ion trap qubits [27].

We also placed the idea of a q-sample into the broader context of an analogy between quantum states and classical probability distributions. If a qpdf can be found that is pure, can be q-sampled, and allows q-stochastic




FIG. 3: a) Quantum Bayesian inference on Bayes net  $\mathcal{B}$  for evidence  $\mathcal{E} = e$  is done by repetition of the circuit shown, with  $k$  incrementing  $k = 0, 1, \dots$ , stopping when the measurement result  $x$  contains evidence bits  $e$ . Then  $x$  can be recorded as a sample from the conditional distribution  $P(Q|\mathcal{E})$ . This corresponds to Algorithm 1. b) The constituents of the Grover iterate  $\hat{G}$ , the state preparation  $\hat{A}_\mathcal{B}$  and phase flip operators  $\hat{S}_e$  and  $\hat{S}_0$ . The state preparation operator is constructed from Theorem 2, and an example is shown in Fig. 1b. The phase flip operators are constructed as shown in Fig. 2.

updates, the quantum machine learning subfield would greatly benefit. Algorithms for many important routines,

such as Metropolis-Hastings, Gibbs-sampling, and even Bayesian learning, could find square-root speedups in a similar manner to our results here.

Artificial intelligence and machine learning tasks are often at least NP-hard. Although exponential speedups on such problems are precluded by BBBV [28], one might hope for square-root speedups, as we have found here, for a variety of tasks. For instance, a common machine learning environment is online or interactive, in which the agent must learn while making decisions. Good algorithms in this case must balance exploration, finding new knowledge, with exploitation, making the best of what is already known. The use of Grover's algorithm in reinforcement learning has been explored [29], but much remains to be investigated. One complication is that machine learning often takes place in a classical world; a robot is not usually allowed to execute a superposition of actions. One might instead focus on learning tasks that take place in a purely quantum setting. For instance, quantum error correcting codes implicitly gather information on what error occurred in order to correct it. Feeding this information back into the circuit, could create an adaptive, intelligent error correcting code.

- 
- [1] S. J. Russell, P. Norvig, *Artificial Intelligence: A Modern Approach*, 2nd Edition, Pearson Education, 2003.
  - [2] M. Bensi, A. D. Kiureghian, D. Straub, Efficient bayesian network modeling of systems, *Reliability Engineering and System Safety* 112 (0) (2013) 200 – 213.
  - [3] R. E. Neapolitan, *Learning Bayesian Networks*, illustrated edition Edition, Prentice Hall, 2003.
  - [4] G. Cooper, E. Herskovits, A bayesian method for the induction of probabilistic networks from data 9 (4) (1992) 309–347.
  - [5] N. Friedman, M. Goldszmidt, A. Wyner, Data analysis with bayesian networks: A bootstrap approach (1999).
  - [6] F. V. Jensen, *Bayesian Networks and Decision Graphs* (Information Science and Statistics), Springer, 2001.
  - [7] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, E. Teller, Equation of state calculations by fast computing machines, *The Journal of Chemical Physics* 21 (6) (1953) 1087–1092.
  - [8] S. Chib, E. Greenberg, Understanding the Metropolis-Hastings Algorithm, *The American Statistician* 49 (4) (1995) 327–335.
  - [9] P. Dagum, M. Luby, Approximating probabilistic inference in bayesian belief networks is np-hard, *Artificial Intelligence* 60 (1) (1993) 141 – 153.
  - [10] V. K. Mansinghka, *Natively Probabilistic Computation*, Ph.D. thesis, MIT (2009).
  - [11] E. Bernstein, U. Vazirani, Quantum complexity theory, in: *Proc. 25th Annual ACM Symposium on Theory of Computing*, ACM, 1993, pp. 11–20.
  - [12] S. Aaronson, Bqp and the polynomial hierarchy, in: *Proceedings of the Forty-second ACM Symposium on Theory of Computing*, STOC '10, ACM, New York, NY, USA, 2010, pp. 141–150.
  - [13] P. W. Shor, Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer, *SIAM J. on Computing* (1997) 1484–1509.
  - [14] L. K. Grover, A fast quantum mechanical algorithm for database search, in: *ANNUAL ACM SYMPOSIUM ON THEORY OF COMPUTING*, ACM, 1996, pp. 212–219.
  - [15] A. Galindo, M. A. Martin-Delgado, Information and computation: Classical and quantum aspects, *Rev. Mod. Phys.* 74 (2002) 347–423.
  - [16] M. A. Nielsen, I. L. Chuang, *Quantum Computation and Quantum Information*, 1st Edition, Cambridge University Press, 2004.
  - [17] S. Jordan, Quantum algorithm zoo, <http://math.nist.gov/quantum/zoo/>.
  - [18] M. Ozols, M. Roetteler, J. Roland, Quantum rejection sampling, in: *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, ITCS '12, ACM, New York, NY, USA, 2012, pp. 290–308.
  - [19] L. Grover, Rapid sampling through quantum computing, *STOC '00*.
  - [20] A. Bookatz, Qma-complete problems arXiv:1212.6312.
  - [21] G. Brassard, P. Høyer, M. Mosca, Quantum amplitude amplification and estimation, 2002, pp. 53–74.
  - [22] D. Aharonov, A. Ta-Shma, Adiabatic quantum state generation and statistical zero knowledge, in: *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, STOC '03, ACM, New York, NY, USA, 2003, pp. 20–29.
  - [23] V. Bergholm, J. J. Vartiainen, M. Möttönen, M. M. Salomaa, Quantum circuits with uniformly controlled one-qubit gates, *Phys. Rev. A* 71 (2005) 052330.
  - [24] D. M. Chickering, D. Heckerman, C. Meek, Large-sample learning of bayesian networks is np-hard, *J. Mach. Learn.*

- Res. 5 (2004) 1287–1330.
- [25] M. Druzdzel, H. Simon, Causality in bayesian belief networks, in: In Proceedings of the Ninth Annual Conference on Uncertainty in Artificial Intelligence (UAI-93, Morgan Kaufmann Publishers, Inc, 1993, pp. 3–11.
  - [26] M. Saeedi, M. Pedram, Linear-depth quantum circuits for n-qubit toffoli gates with no ancilla, arXiv preprint arXiv:1303.3557.
  - [27] D. Hanneke, J. P. Home, J. D. Jost, J. M. Amini, D. Leibfried, D. J. Wineland, Realization of a programmable two-qubit quantum processor, Nat Phys 6 (1) (2010) 13–16.
  - [28] C. H. Bennett, E. Bernstein, G. Brassard, U. Vazirani, Strengths and weaknesses of quantum computing, SIAM JOURNAL OF COMPUTATION (1997) 1510–1523.
  - [29] D. Dong, C. Chen, H. Li, T. Tarn, Quantum reinforcement learning, Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on 38 (5) (2008) 1207–1220.